# Breaking the cycle of complexity

Danica Porobic

Oracle

#### Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing and price of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle corporation. Fees apply for new database product offerings.





© Matt Turck (@mattturck) and ShivonZilis (@shivonz)

## Multisocket multicores



#### **Challenge: non-uniform communication**

## **OLTP on Hardware Islands**



Shared-everything

stablenot optimal

Island shared-nothing

Shared-nothing

- robust middle ground
- 🗸 fast
- sensitive to workload

#### **Optimal configuration depends on workload and hw**

## Rack-scale hardware platforms



- Abundant non-uniform parallelism
  Need to scale across many cores
- Large main memories
  - Datasets are memory resident
- Network & DRAM converge
  - Need to scale across multiple nodes

### A hierarchy of Hardware Islands



#### **Communication costs determine throughput**



## Why don't updates scale out?

2nodes x 1core





#### **Multicore-optimized OCC very sensitive to delays** 10

## **DBMS** Layer-Cake

Queries		Query Parser	Transactions
		Query Optimization and Execution	
	Concurrency Control Recovery	Relational Operators	
		Buffer Management	
		Files and Access Methods	
		Disk Space Management	

## **DBMS Layer-Cake**

	Query Parser	Transactions
	Query Optimization and Execution	
	Relational Operators	
Concurrency Control	Buffer Management	
	Files and Access Methods	
Recovery	Disk Space Management	

## **DBMS** Layer-Cake

Queries	Query Parser				
	Query Optimization and Execution				
	Relational Operators				
Concurrency	Buffer Management				
	Files and Access Methods				
Recovery	Disk Space Management				











#### **BIG DATA LANDSCAPE 2017**



st updated 5/3/2017

EARLY STAGE VENTURE CAPITAL





















## Defeat complexity!

Big data and new applications pose new challenges

- Deluge of specialized systems
- All-to-all connectors/ETL tools

- Rethink specialization
  - Flexible functional components
  - Composable like toy blocks

